Geodesics using Waves: Computing Distances using Wave Propagation

Ayushi Sinha* Michael Kazhdan†
Department of Computer Science, the Johns Hopkins University, Baltimore, MD

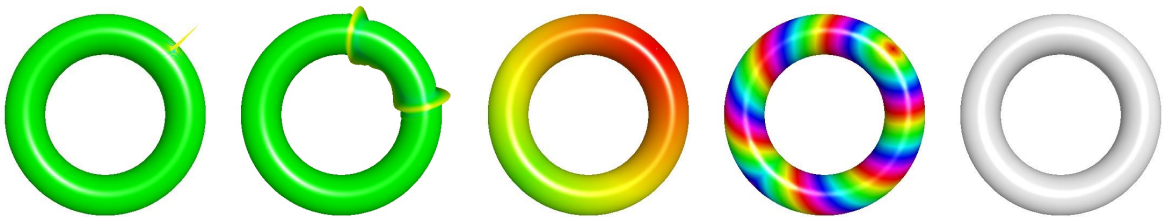

Figure 1: *Approximate geodesic distances using wave propagation: (a) Wave propagation is initiated by setting the initial signal such that its only peak is at the source of the wave; (b) repeatedly solving the wave equation on this signal propagates the wave outward from the source, across the shape; here yellow indicates the positive part of the wavefront, and green indicates that the height of the wave is approximately zero; (c) the wave propagation is used to define a pseudo-distance function, (d) which, in turn, is used to approximate geodesic distances; (e) the error compared to exact geodesics, where 0 error is indicated by white.*

arXiv:1612.02509v1 [cs.CG] 8 Dec 2016

Abstract

In this paper, we present a new method for computing approximate geodesic distances. We introduce the wave method for approximating geodesic distances from a point on a manifold mesh. Our method involves the solution of two linear systems of equations. One system of equations is solved repeatedly to propagate the wave on the entire mesh, and one system is solved once after wave propagation is complete in order to compute the approximate geodesic distances up to an additive constant. However, these systems need to be pre-factored only once, and can be solved efficiently at each iteration. All of our tests required approximately between 300 and 400 iterations, which were completed in a few seconds. Therefore, this method can approximate geodesic distances quickly, and the approximation is highly accurate.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: geodesic distance, wave propagation, Poisson equation

1 Introduction

When a pebble is dropped in a body of water, the impact causes waves to form and propagate in all directions moving away from the point of impact. If the medium is uniform, then the waves will travel uniformly in every direction, and therefore reach points that are equidistant from the source of the wave at the same time. Similarly, waves can also be propagated on manifold meshes, and expected to reach points equidistant from the source of the wave at approximately the same time. Therefore, we can define a function that measures the time required for a wave to travel from the source of the wave, x, to any other point, y, on the mesh.This function can be thought of as a pseudo-distance function, since it does not tell us what the real distance between x and y is, but it gives us some sense of how far the two points are from each other.

This pseudo-distance function can be used to approximate exact geodesic distances. We compute the normalized gradient of the pseudo-distance function for every face in the mesh, and then solve a Poisson equation to compute a function which has the same gradient as the normalized gradient of the pseudo-distance function. This produces the approximate geodesic distances we want, up to an additive constant [Crane et al. 2013].

2 Related Work

Solving the special case of the Eikonal equation

$$|\nabla \phi| = 1 \tag{1}$$

subject to boundary conditions $\phi|_{\partial\Omega} = 0$, produces the shortest distance from the boundary, $\partial\Omega$, to any point in Ω, or the geodesic distance. However, this partial differential equation is non-linear and hard to solve without using some type of iterative relaxation approach [Hysing and Turek 2005]. The fast marching method [Kimmel and Sethian 1998] and fast sweeping method [**?**] on triangulated manifolds are two popular algorithms for computing geodesic distances using this approach. [Mitchell et al. 1987] developed another method to compute the exact piecewise linear geodesic distance between a source and all other vertices of a manifold mesh. However, this algorithm was not implemented until much later by [Surazhsky et al. 2005]. They also extended the work of [Mitchell et al. 1987] with an approximation algorithm with bounded errors and faster runtime.

One of the most recent methods for approximating geodesics is the heat method [Crane et al. 2013], which offers several improvements over previous methods. One of the advantages it offers is that it is simple to implement since it requires the solution of two standard linear systems of equations which can be solved efficiently, making this method very fast. Our method is similar to this method in that our method also requires the solution of two linear systems of equations, which can be prefactored [Chen et al. 2008], and then solved efficiently. However, our method requires one of these systems of equations to be solved repeatedly.

*e-mail:asinha8@jhu.edu
†email:misha@cs.jhu.edu

3 The Wave Method

Let $\mathcal{M} = (\mathcal{V}, \mathcal{E}, \mathcal{F})$ be a manifold mesh with $\mathcal{V}$ vertices, $\mathcal{E}$ edges, and $\mathcal{F}$ faces. Then, the wave equation defined over $\mathcal{M}$ can be written as

$$\begin{aligned} \phi_t : \mathcal{M} \times \mathbb{R} &\to \mathbb{R} \\ \frac{\partial^2}{\partial t^2}\phi_t &= \mu \Delta \phi_t, \\ \phi_0 &= \phi \\ \frac{\partial}{\partial t}\phi_0 &= 0 \\ \mu &= 1 \end{aligned} \tag{2}$$

where Δ is the Laplace operator, and ϕ is the initial signal. We will talk about how we set the initial signal in the next section. We solve this equation to propagate a wave over time on the surface of $\mathcal{M}$.

By tracking this wave propagation, we can define a pseudo-distance function over the manifold. We can then compute the normalized gradient, G, of the pseudo-distance function, and solve a Poisson equation with fixed boundary to approximate a function with the same gradient as G. Since, the gradient of the geodesic distance function has unit length, the solution of the Poisson equation gives us the approximate geodesic distances over the manifold up to an additive scale [Crane et al. 2013].

3.1 Wave Propagation

We start the wave propagation at a source vertex, v, by setting an initial signal, ϕ_0, such that its only peak is at v. Then, we flow Eq. 2 repeatedly, and track the leading wavefront until we have recorded a pseudo-distance at each vertex in the mesh. We will describe our initial signal and wave propagation in this section.

3.1.1 Initial Value

The initial signal is set so that the signal is high only at the vertex that is chosen to be the source of the wave. The signal at all other vertices is set to 0, while at the source vertex it can be set to any value greater than zero. We set the initial signal such that for each vertex, $v \in \mathcal{V}$,

$$\phi_0(v) = \begin{cases} 1, & \text{if } v \text{ is the source of the wave,} \\ 0, & \text{otherwise.} \end{cases}$$

Since, the change in this signal at time $t = 0$ is defined to be 0 (Eq. 2), the signal in the time-step before the first time-step is defined to be the same as ϕ_0.

3.1.2 Implicit Formulation

We discretize the flow using the following implicit formulation for a mesh with n vertices:

$$\left\langle \frac{\phi_{t+\delta} - 2\phi_t + \phi_{t-\delta}}{\delta^2}, b_v \right\rangle = \langle \Delta \phi_{t+\delta}, b_v \rangle ,$$

where $b_v, v = 1 \dots n$ is the hat-basis function at vertex $v \in \mathcal{V}$, and δ is the time-step. By expressing ϕ_t as a linear combination of these basis functions,

$$\phi_t(p) = \sum_{w \in V} \xi_w(t) b_w(p),$$

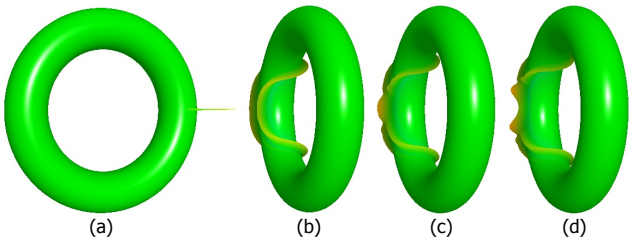

Figure 2: *When two waves with equal amplitudes traveling towards each other collide, the amplitude of the wave at collision is the sum of the amplitudes of the two colliding waves. This is called constructive collision. We can see an instance of constructive collision on the torus model. (a) shows the source of the wave, and (b) shows the wavefront just before collision. Constructive collision resulting in a wave with a higher amplitude than the colliding wavefronts can be seen as the bump where the collision occurred in (c). (d) shows these bumps propagating past the point of collision.*

we can write

$$\begin{aligned} &\frac{1}{\delta^2} \sum_{w \in V} \langle \xi_w(t+\delta) b_w - 2\xi_w(t) b_w + \xi_w(t-\delta) b_w, b_v \rangle \\ &\qquad = \sum_{w \in V} \langle \Delta \xi_w(t+\delta) b_w, b_v \rangle , \forall v \in V \\ \Rightarrow &\frac{1}{\delta^2} \sum_{w \in V} \left[\xi_w(t+\delta) - 2\xi_w(t) + \xi_w(t-\delta) \right] \langle b_w, b_v \rangle \\ &\qquad = \sum_{w \in V} \xi_w(t+\delta) \langle \Delta b_w, b_v \rangle , \forall v \in V \end{aligned} \tag{3}$$

Since, $M = \langle b_w, b_v \rangle$ is the mass matrix, and $S = -\langle \Delta b_w, b_v \rangle$ is the stiffness matrix, Eq. 3 becomes

$$\left[2\vec{\xi}(t) - \vec{\xi}(t-\delta) \right] M = (\mu \delta^2 S + M) \vec{\xi}(t+\delta) \tag{4}$$

This equation is of the form $Ax = b$, where x is unknown, and A is a symmetric, positive definite matrix.

The source of the wave and the wave propagation produced can be seen in Figures 2a and 2b, respectively. We want to be able to track the leading tall *wavefront* traveling from the source across the shape in order to define a pseudo-distance function on $\mathcal{M}$.

3.2 Pseudo-Distance Function

3.2.1 Naive Approach

Since we expect the largest amplitude of the wave to be at the wavefront, we first tried to track it by tracking its peak. That is, for each vertex i on the manifold mesh, we record the time at which the wave attains its maximum value. This distance function, d, can be formulated as

$$\begin{aligned} d : \mathcal{V} &\to \mathbb{R} \\ d(v) &= \arg\max_t \phi_t(v). \end{aligned}$$

This distance formulation, however, runs into problems at points of collision. When there is a constructive collision, the amplitude of the wave at the point of collision is the sum of the amplitudes of the two colliding waves. Therefore, at the point of collision, the amplitude of the wave becomes larger than the amplitude of the wavefront (Fig. 2c,d), causing artifacts in the distance function.

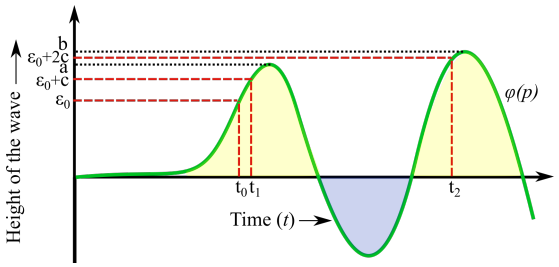

Figure 3: *Instability of the distance function $d(p)$: Given that at $\epsilon = \epsilon_0$, $d(p) = t_0$, adding a small constant c to ϵ_0, such that $\epsilon_0 + c < a$, where a is the amplitude of the first wave in the figure, results in a small change in the distance function, $d(p) = t_1$, which is close to t_0. However, if c is added to ϵ_0 again, such that $\epsilon_0 + 2c > a$, then the change in the distance function is extremely large, since $d(p) = t_2$ is a lot larger than t_1. This error can happen if ϵ is larger than the height of the wavefront, and there is a larger wave following this wavefront due to a collision.*

3.2.2 Our Approach

Instead, we track the time at which the wave first attains a prescribed height. That is, for each vertex v on the mesh, we record the time at which the wave first achieves a small positive value, ϵ.

$$d : \mathcal{V} \to \mathbb{R}$$
$$d(v, \epsilon) = inf\{t : \phi_t(v) = \epsilon\}. \quad (5)$$

This distance formulation records the time at which the wavefront reaches a vertex v. This formulation is able to avoid the problems due to collision that arose in the previous formulation. However, this formulation is unstable.

3.2.3 Instability

As the wavefront propagates across the shape, its amplitude decreases gradually as the energy of the wave dissipates to larger areas and over more vertices. We have to take this into consideration when choosing a value for ϵ, because if the height of the wavefront is smaller than ϵ at a vertex v, then no $d(v, \epsilon)$ value will be recorded for the wavefront. However, a $d(v, \epsilon)$ can be recorded at a later time due to a larger wave resulting from a collision, for instance. In other words, a slight change in ϵ can result in a big change in d (Fig. 3), making it important for ϵ to remain below the expected height of the wavefront throughout the wave propagation. Therefore, we need to choose ϵ conservatively, while also making sure to not choose too small a value, since this will make d sensitive to noise.

3.2.4 Spatially Varying ϵ

In order to address this problem, we let ϵ vary as we move away from the source. The idea is to estimate the height of the wavefront, $h(i)$, at every iteration, i, of the wave propagation. Using this information, we can estimate a function for ϵ, $\epsilon(i)$, which is smaller than $h(i)$ at every iteration. In order to do this, we recorded the maximum height of the wave with $\delta = 0.05$ at each iteration on a sphere model to avoid peaks resulting from collision (Fig. 4a):

$$h(i) = \arg\max_v \phi(v), \text{ for } i = 1, 2, \ldots$$

Empirically, we observed that the height of the wave was approximately inversely proportional to the iteration of wave propagation. From the graph, we expect $h(i) \approx ci^a$, where i is the iteration, c is a constant, and a is a negative number. Assuming $c = 1$, since it is easier to estimate than a, we plot $\ln h(i)$ against $\ln i$ (Fig. 4b) to estimate the slope, a, of the graph:

$$h(i) = i^a \Rightarrow \ln h(i) = a \ln i$$
$$\Rightarrow \frac{\ln h(i)}{\ln i} = a$$

We found the average slope for $\delta = 0.05$ to be approximately -3. Therefore,

$$h(i) = \frac{c}{i^3}$$

Now, we can compute the constant c. We want $\epsilon(i)$ to vary in the same manner as $h(i)$. At the same time, we also want $\epsilon(i)$ to remain lower than $h(i)$ at every iteration. Therefore, we set c to be equal to half the maximum height of the wave at the first iteration (Fig. 5a).

$$\epsilon(i) = \frac{h(1)}{2i^3}$$

We also observed wave propagation with varying δ to understand how the height of the wave changes as it propagates, and found the value of a to vary with δ. As δ increases, the rate at which the height of the wave falls per iteration also increases. Therefore, a decreases as δ increases (Fig. 5b).

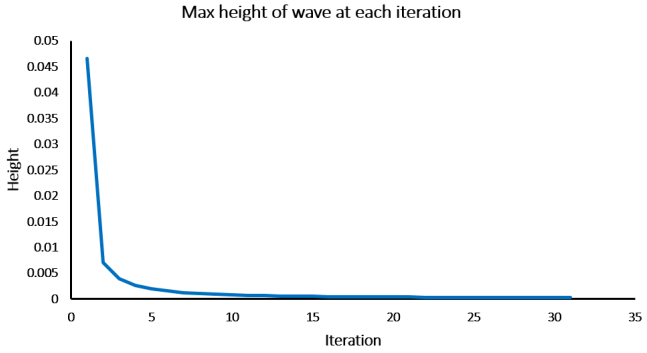

(a)

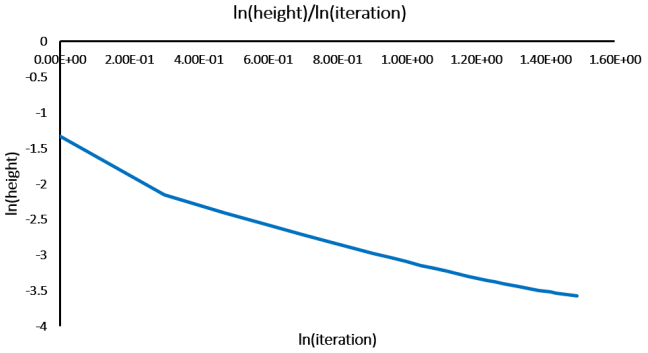

(b)

Figure 4: *Wave propagated on a sphere model with $\delta = 0.05$. (a) Maximum height of the wave plotted per iteration; (b) the plot for log(height)/log(iteration) has average slope of -3.*

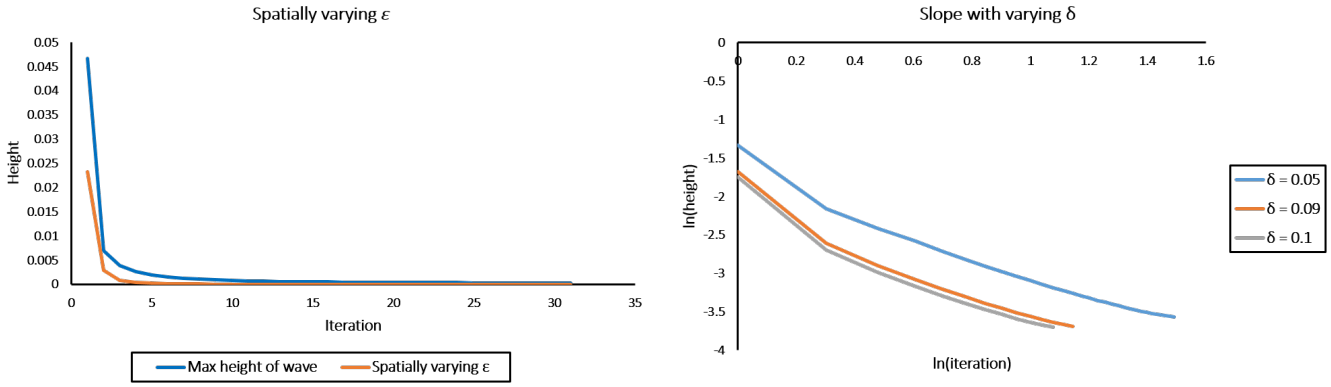

Figure 5: *(a) Spatially varying ϵ plotted per iteration along with the maximum height of the wave; (b) plot for log(height)/log(iteration) for varying δ.*

3.2.5 Time Estimation

Now that we have a formulation for ϵ for every iteration, we need to compute $d(v, \epsilon)$ at all the vertices as shown in Eq. 5. That is, we need to record the time at which $\phi_t(v)$ becomes exactly equal to the ϵ at that iteration. Since absolute equality is improbable, we need to interpolate to estimate when the wave achieved height $= \epsilon$. We linearly interpolate between the time $t1$ at which height $< \epsilon$ was achieved and $t2$ at which height $> \epsilon$ was achieved to estimate the time t at which height $= \epsilon$.

3.3 Geodesic Distance Function

Once we have a pseudo-distance function defined at every vertex, we can compute the normalized gradient, G, of this function, and use it to find a function with the same gradient as G by solving a Poisson equation as in [Crane et al. 2013].

Since the gradient of the true distance function has unit length, the function we obtain by solving the Poisson equation is an approximation to the geodesic distance function we want, up to an additive constant.

3.3.1 Gradient

For a manifold $\mathcal{M} = (\mathcal{V}, \mathcal{E}, \mathcal{F})$, we can compute the gradient, G, of our pseudo-distance function defined on the vertices ($\mathcal{V}$) of our mesh:

$$G : \mathbb{R}^{|\mathcal{V}|} \rightarrow \mathbb{R}^{2|\mathcal{T}|}$$

We compute normalized G, $\bar{G}$, per mesh face $\mathcal{F}$, and compute its divergence. We can compute two gradient fields per mesh face:

$$\begin{aligned} \vec{w_1} &= p(v_1 - v_0) + q(v_2 - v_0) \\ \vec{w_2} &= s(v_1 - v_0) + t(v_2 - v_0) \end{aligned} \tag{6}$$

where p, q, s and t are unknown. Then,

$$\begin{aligned} \langle \vec{w_1}, \vec{w_2} \rangle &= \langle p(v_1 - v_0) + q(v_2 - v_0), s(v_1 - v_0) + t(v_2 - v_0) \rangle \\ &= \begin{pmatrix} s & t \end{pmatrix} K \begin{pmatrix} p \\ q \end{pmatrix} \end{aligned}$$

where $K_{ij} = \langle v_i - v_0, v_j - v_0 \rangle, i, j = 1, 2$.

We also know by definition that the gradients satisfy

$$\begin{aligned} \langle \vec{w_1}, v_1 - v_0 \rangle &= d_1 - d_0 \\ \langle \vec{w_1}, v_2 - v_0 \rangle &= d_2 - d_0 \end{aligned} \tag{7}$$

$$\begin{aligned} \langle \vec{w_2}, v_1 - v_0 \rangle &= d_1 - d_0 \\ \langle \vec{w_2}, v_2 - v_0 \rangle &= d_2 - d_0 \end{aligned} \tag{8}$$

where d_v are values of the pseudo-distance function at each vertex. We can rewrite Eq. 7 as

$$\begin{aligned} \langle p(v_1 - v_0) + q(v_2 - v_0), v_1 - v_0 \rangle &= d_1 - d_0, \\ \langle p(v_1 - v_0) + q(v_2 - v_0), v_2 - v_0 \rangle &= d_2 - d_0. \end{aligned}$$

We have two equations and two unknowns, p and q, which can be computed by solving

$$K \begin{pmatrix} p \\ q \end{pmatrix} = \begin{pmatrix} d_1 - d_0 \\ d_2 - d_0 \end{pmatrix}$$

Similarly, we can compute s and t, which allows us to compute the normalized Eq. 6 per mesh face. Then, we define a difference function at each edge in the mesh by computing the mean gradient at the two triangles incident upon each edge:

$$\begin{aligned} \widetilde{G} &: \mathbb{R}^{|\mathcal{T}|} \rightarrow \mathbb{R}^{|\mathcal{E}|} \\ e = (v_i, v_j) &\rightarrow \frac{1}{2} \langle v_j - v_i, w_{ij}^1 + w_{ij}^2 \rangle, \end{aligned} \tag{9}$$

where $\mathcal{E}$ is the set of edges.

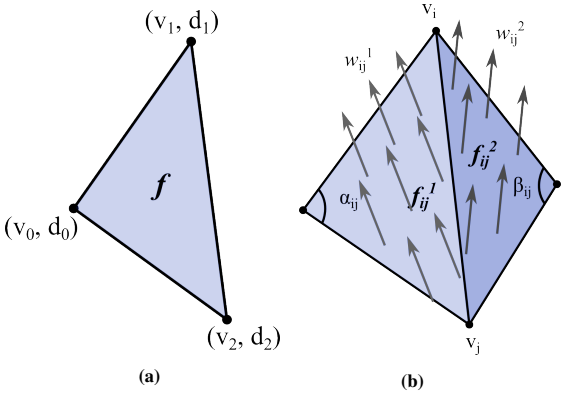

Figure 6

Table 1: *This table shows a numerical comparison between the wave method, heat method, and the exact method.*

MODEL	FACES	WAVE METHOD					HEAT METHOD				EXACT
		δ	MEAN RAW ERROR	MEAN RELATIVE ERROR	MAX RAW ERROR	TIME (s)	MEAN RAW ERROR	MEAN RELATIVE ERROR	MAX RAW ERROR	TIME (s)	TIME (s)
BUNNY	28k	0.005	0.00574625	0.00500049	0.058528	2.5	0.0137430	0.0112878	0.0974171	0.205	0.573
HORSE	96k	0.006	0.0156859	0.0131532	0.0560383	6.288	0.00470463	0.00425243	0.0312462	0.859	4.016
KITTEN	106k	0.005	0.00684904	0.00505089	0.0367825	8.101	0.00478574	0.00453304	0.0197093	1.477	5.335
BIMBA	149k	0.004	0.00693687	0.00742395	0.06061	12.579	0.00773516	0.00741442	0.0512189	2.476	7.641
RAMESSES	1.6M	0.09	0.0245474	0.0255807	0.131039	7.175	0.0111101	0.00909884	0.237803	36.227	37.638

3.3.2 Poisson Equation

Once we have gradients defined at each edge, we solve the Poisson equation to obtain the geodesic distance function on the mesh. Let L be the Laplacian which can be computed as

$$L_{ij} \in \mathbb{R}^{|\mathcal{V}|\times|\mathcal{V}|} = \begin{cases} \frac{\cot\alpha_{ij}+\cot\beta_{ij}}{2}, & \text{if } j \in N(i) \\ -\sum_{k\in N(i)} L_{ik}, & \text{if } i = j \\ 0, & \text{otherwise,} \end{cases}$$

where α_{ij} and β_{ij} are as shown in Fig. 6b, and $N(i)$ is the 1-ring neighborhood of vertex i. L can also be defined as

$$L = D^t \Lambda D,$$

where D is the matrix representing a linear operator in bases $\{m_1, m_2, \ldots, m_n\}$ and $\{o_1, o_2, \ldots, o_n\}$ defined on real vector spaces M and O, making D^t the matrix representing the dual operator in the corresponding dual bases $\{m_1^*, m_2^*, \ldots, m_n^*\}$ and $\{o_1^*, o_2^*, \ldots, o_n^*\}$ in M^* and O^*, and $\Lambda = \langle l, m\rangle$, $l, m \in M$ is the map from real vector space to dual space. We define D and Λ as:

$$D_{ev} \in \mathbb{R}^{|\mathcal{E}|\times|\mathcal{V}|} = \begin{cases} 1, & \text{if vertex } v \text{ is the end point of edge } e \\ -1, & \text{if vertex } v \text{ is the starting point of } e \\ 0, & \text{otherwise,} \end{cases}$$

and

$$\Lambda_{ef} \in \mathbb{R}^{|\mathcal{E}|\times|\mathcal{E}|} = \begin{cases} \frac{\cot\alpha_{ij}+\cot\beta_{ij}}{2}, & \text{if } e = f \\ 0, & \text{otherwise.} \end{cases}$$

We want to compute

$$\begin{aligned} & w = (D^t\Lambda D)^{-1} D^t \Lambda \tilde{G} \\ \Rightarrow \quad & w = L^{-1} D^t \Lambda \tilde{G} \\ \Rightarrow \quad & Lw = D^t \Lambda \tilde{G}, \end{aligned}$$

which is an equation of the form $Ax = b$, where $x = w$ is an unknown, and we can compute L, D, Λ and $\tilde{G}$. A is a symmetric, positive definite matrix. We solve this equation for w which is the geodesic distance function we want up to an additive constant. We shift this distance function so that the smallest distance is 0. This gives us our final geodesic distances.

Table 2: *This table shows runtime comparison between the wave method and the exact method for increasing number of faces.*

FACES	WAVE METHOD				EXACT
	MEAN RAW ERROR	MEAN RELATIVE ERROR	MAX ERROR	TIME (s)	TIME (s)
20k	0.016717	0.026019	0.051859	2.031	0.388
80k	0.012966	0.027448	0.040692	5.32	3.682
320k	0.007401	0.01482	0.020993	18.198	44.233

4 Evaluation and Results

We compare our approximate geodesics using wave propagation against geodesics from heat [Crane et al. 2013], as well as against exact geodesics, described in [Surazhsky et al. 2005] and implemented by Kirsanov [Kirsanov et al] (Table 1). We compute two errors to assess the accuracy of our method:

- Raw error: the absolute difference between geodesics computed using our method and the exact method.
- Relative error: the ratio between the raw error and the exact geodesic distance.

These metrics are computed at each vertex and then averaged. We also visually compare our method with the exact method in Fig. 8. We indicate zero error by an RGB value of 1, and an error of 1 by an RGB value of 0 (Fig. 8, rows 3, 5), since all our errors are smaller than 1 (Table 1). However, these errors are hard to visualize because they are much smaller than 1, and, in addition, shadows due to the lighting model also manifest themselves in similar colors as the errors. Therefore, we visualize the models without lighting, and rescale the intensity values so that values between 0.7 and 1 are remapped to values between 0 and 1 (Fig. 8, rows 4, 6). We see that raw error is higher away from the source, since the wave becomes more smoothed as it travels away from the source. Some more results using our method can be seen in Fig. 14.

Both our method and the heat method solve two linear systems. The heat method requires that both these systems be solved only once, whereas our method requires us to solve one of the two linear systems repeatedly until our propagating wave has traveled across the shape. This makes our method slower than the heat method. How-

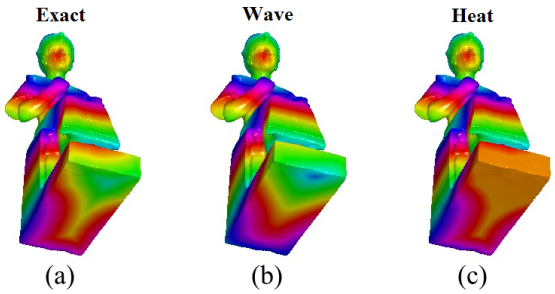

Figure 7: *The Ramesses model shows that our method can handle very large meshes (1.6 million faces), and approximate geodesics reliably even far away from the source. The isocontours on the base of the Ramesses model plotted using geodesics from the wave method (b) show a more similar coloring compared with exact geodesics (a), than those plotted using heat geodesics (c). However, the shape of the isocontours using our method (b) is more smoothed than both the exact method (a) and the heat method (c).*

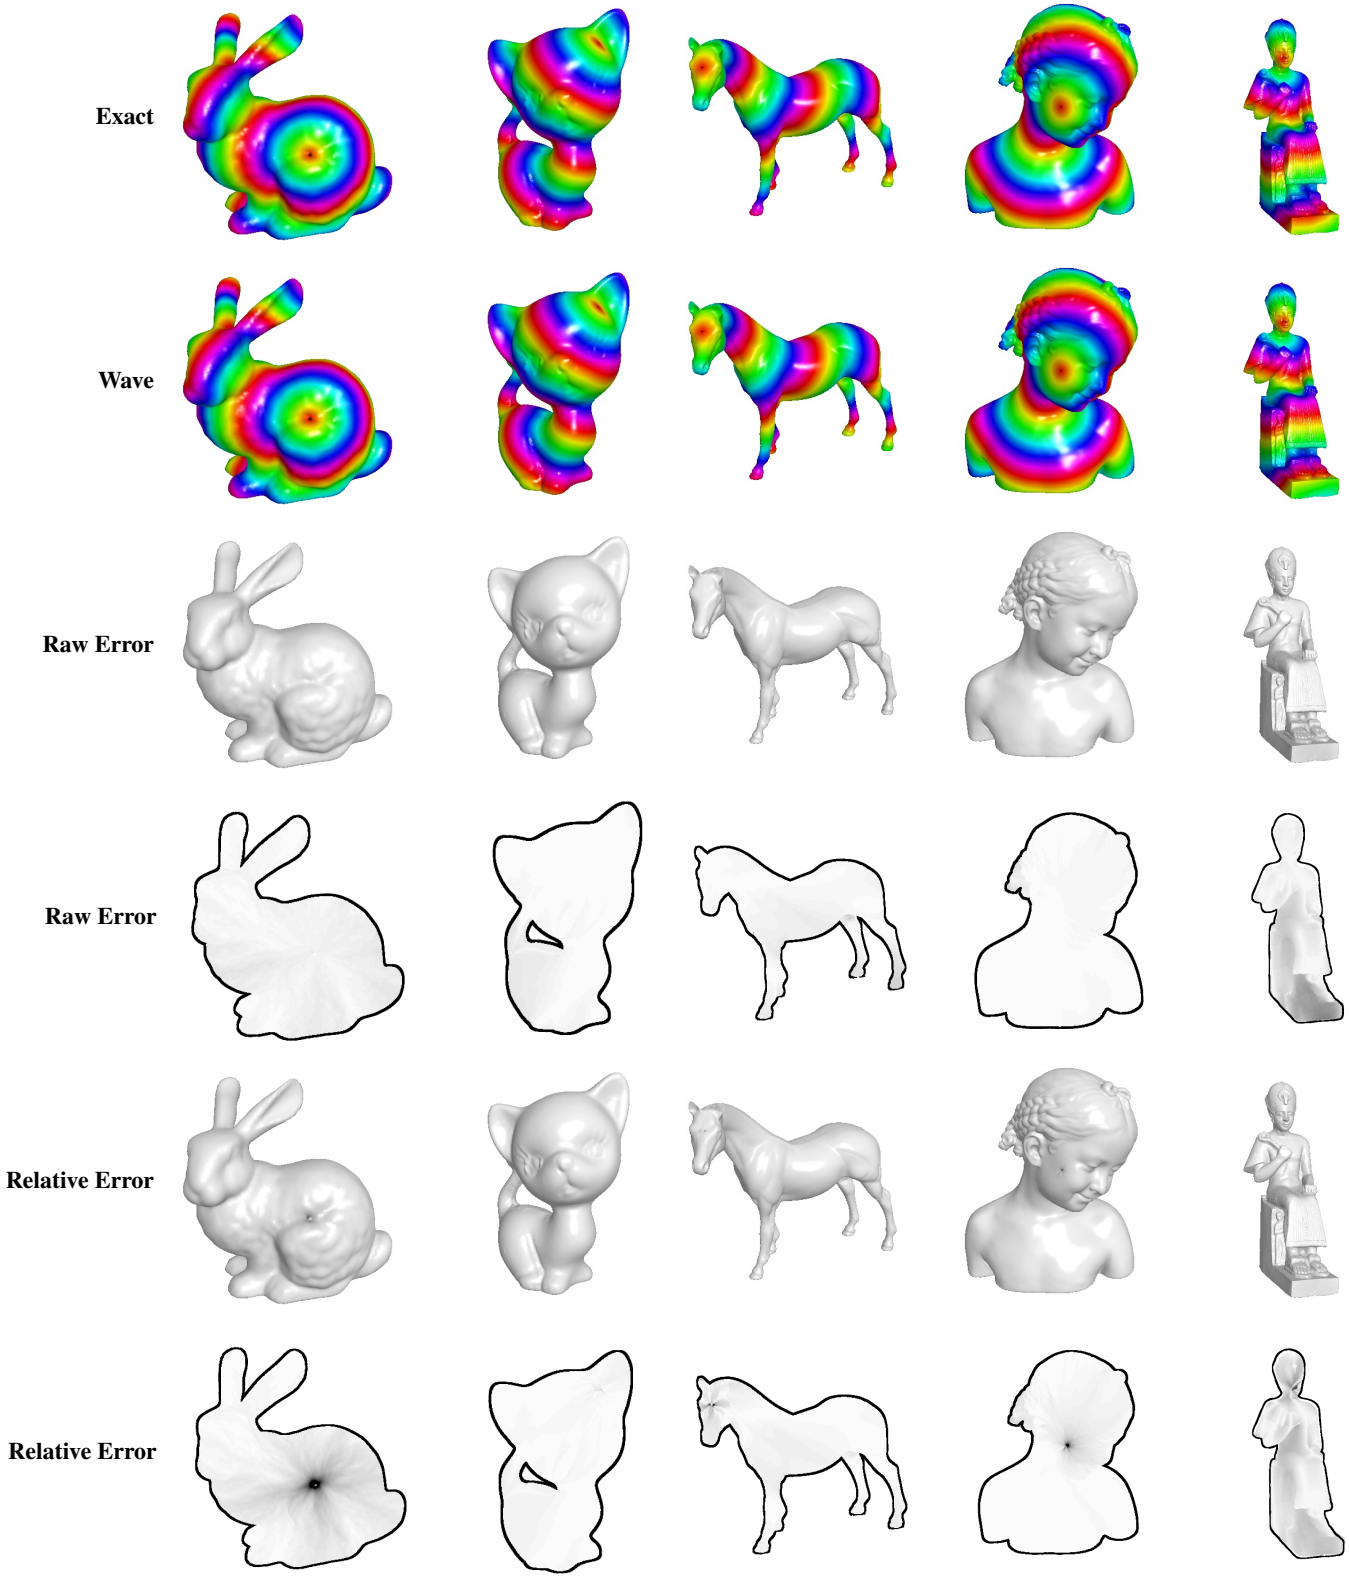

Figure 8: *Comparison of the wave method with the exact geodesic method on (left to right) the Bunny, Kitten, Horse, Bimba, and Ramesses models. This figure shows isocontours from exact geodesics (top row), isocontours from geodesics obtained using our method (2nd row), raw error (3rd row), raw error without lighting (4th row), relative error (5th row), and relative error without lighting (last row). The images showing error without lighting are rescaled to show where the errors are located. The black dot in the 1st and 2nd rows indicates the source.*

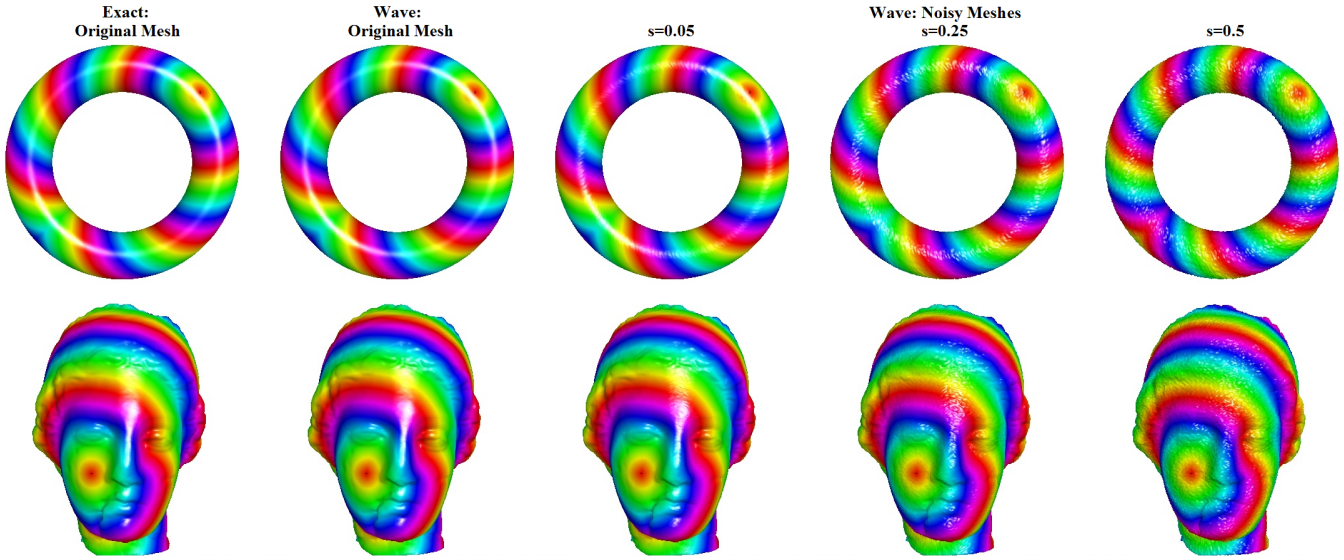

Figure 9: *Geodesic distance computed on the torus (top) and Igea (bottom) models. For both shapes, (from left to right) we have the exact geodesics and geodesics from waves on the original models, as well as geodesics from waves on the models with added noise, where s = 0.05, 0.25 and 0.5.*

ever, our method can be sped up by changing the parameters that increase the speed of the wave traveling on the surface of our mesh. The runtime of the wave method is affected by parameter δ. The larger these parameters are, the faster the wave propagates. However, errors in geodesics approximated using our method increase as the time-step, or δ, increases (Fig. 10a). The potential for aliasing also increases as δ gets larger. On the other hand, increasing the resolution of a mesh increases the accuracy of our method, while also increasing the time required to compute the geodesics. However, we found that our method scales better than the exact method (Table. 2).

We further evaluated the robustness of our method in the presence of noise [Rusinkiewicz 2004]. We added noise to each vertex by shifting each vertex by

$$O(s \times l),$$

where l is the median edge length of all the edges incident on vertex v, and s varies between 0 and 1. Our method remained fairly robust in the presence of noise (Fig. 9), producing good approximations even in the presence of considerable amounts of noise (Fig. 10b). We also compared errors on smoothed as well as sharpened versions of models [Rusinkiewicz 2004]. We smoothed our models by performing m iterations of simple umbrella smoothing [Kobbelt et al. 1998], where m ranges between 10 and 50. In order to sharpen our models, we first smooth our model using a smoothing kernel with radius equal to σ, which is given by

$$\sigma = s \times l,$$

where l is the median edge length of all the edges incident on vertex v. In our experiments, s varies between 0.005 and 5. Once the position of each vertex after smoothing is obtained, we add to the original vertices twice the difference between the original and the new vertices to obtain a sharpened mesh. Our method is quite robust to smoothening and sharpening, with minimal changes in error. As meshes become sharper, the errors tend to increase slightly, whereas the errors decrease as meshes become smoother. The increase in error with increasing sharpness arises from discretization.

Our method also performs well in the presence of holes. The geodesics computed on a mesh without holes, and the same mesh with holes, are comparable, and both have small errors. For instance, on the bunny model without holes (Fig. 11a,b), our approximations had a raw error of 0.0057, and a relative error of 0.005. The errors on the bunny model with holes (Fig. 11c,d) was slightly higher. The raw error for the model with holes was 0.013, and the relative error was 0.012.

5 Conclusion and Future Work

We present a new method to approximate geodesic distances on shapes. Our method produces reliable approximations to exact geodesic distances. We also show that our method scales well with the size of the mesh. Our method can therefore be used in several applications that require approximate geodesics. However, there are several areas that can still be explored. For instance, the method for tracking wave propagation can be improved, and the current height function for ϵ can also be further optimized. Our method is also very sensitive to the value of δ. Making it robust to δ is an important future direction of work.

However, most importantly, our current divergence computation can be improved by storing the gradient of our pseudo-distance function at triangle faces, rather than storing the mean gradient at every edge (Eq. 9). Although we have not implemented this, the details of this implementation are explained in Appendix A.

This method has several applications. One application, for instance, is estimating Voronoi regions by propagating waves from multiple sources. Another application is that it can be used in physically based rendering of subsurface scattering to quickly estimate points that are geodesically close by, instead of computing nearby points using Euclidean distances.

Acknowledgments

This research was supported by the Johns Hopkins University internal funds. We would also like to thank the Stanford Computer Graphics Laboratory and AIM@Shape for providing us with meshes to test our method.

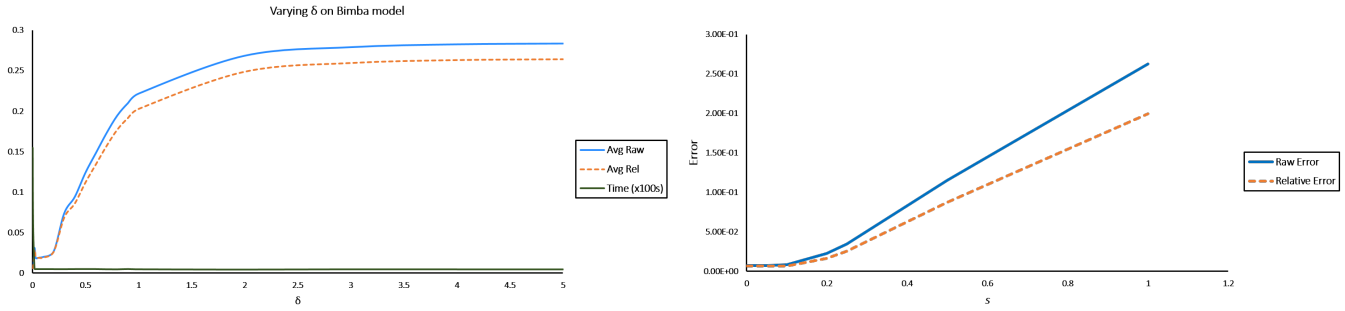

(a) (b)

Figure 10: *(a) This graph shows how the time-step, δ, affects raw and relative errors, and time ($\times 100s$) taken to compute geodesics on the bimba model. As δ increases, errors also increase, while the time taken decays very quickly. For values of δ between 0.005 and 0.1, the errors are small, and the time taken to compute geodesics is under $10s$; (b) Errors computed for increasing noise.*

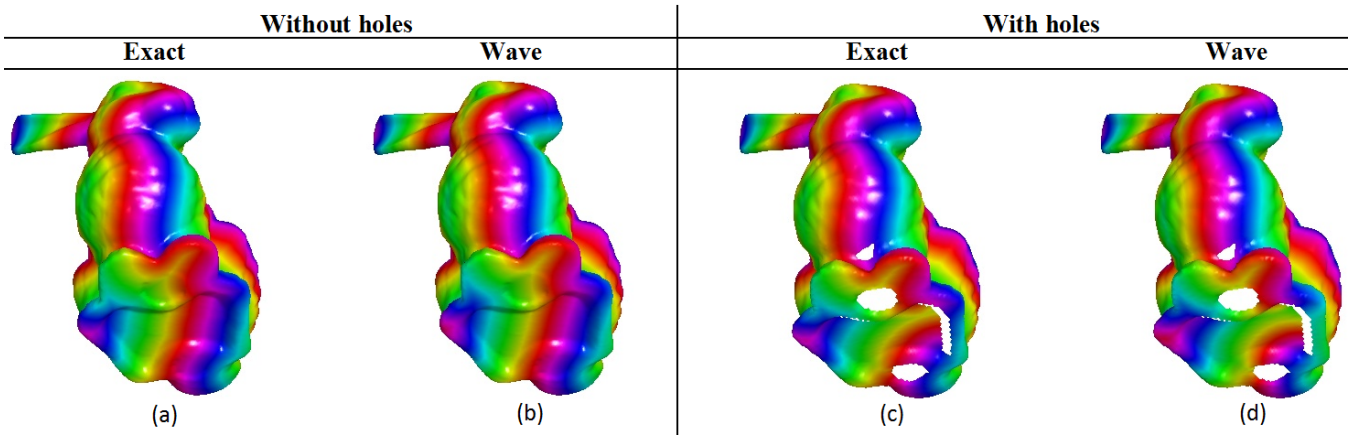

Figure 11: *(a) Exact geodesics, and (b) approximate geodesics obtained using the wave method, plotted on the Stanford bunny model without holes. (c) Exact geodesics, and (d) geodesics using waves, plotted on the bunny model with holes.*

References

CHEN, Y., DAVIS, T., HAGER, W., AND RAJAMANICKAM, S. 2008. Algorithm 887: Cholmod, supernodal sparse cholesky factorization and update/downdate. *ACM Transactions on Mathematical Software (TOMS) 35*, 3, 22:1–22:14.

CRANE, K., WEISCHEDEL, C., AND M, W. 2013. Geodesics in heat: A new approach to computing distance based on heat flow. *ACM Transactions on Graphics (TOG) 32*, 5, 152:1–152:11.

HYSING, S., AND TUREK, S. 2005. The eikonal equation: Numerical efficiency vs. algorithmic complexity on quadrilateral grids. *Proceedings of ALGORITMY*, 22–31.

KIMMEL, R., AND SETHIAN, J. 1998. Computing geodesic paths on manifolds. In *Proceedings of the National Academy of Sciences, USA*, vol. 95, 8431–8435.

KOBBELT, L., CAMPAGNA, S., VORSATZ, J., AND SEIDEL, H. 1998. Interactive multi-resolution modeling on arbitrary meshes. *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, 105–114.

MITCHELL, J., MOUNT, D., AND PAPADIMITRIOU, C. 1987. The Discrete Geodesic Problem. *SIAM Journal on Computing 16*, 4, 647–668.

RUSINKIEWICZ, S. 2004. trimesh2.

SURAZHSKY, V., SURAZHSKY, T., KIRSANOV, D., GORTLER, S., AND HOPPE, H. 2005. Fast exact and approximate geodesics on meshes. *ACM Transactions on Graphics (TOG) 24*, 3, 553–560.

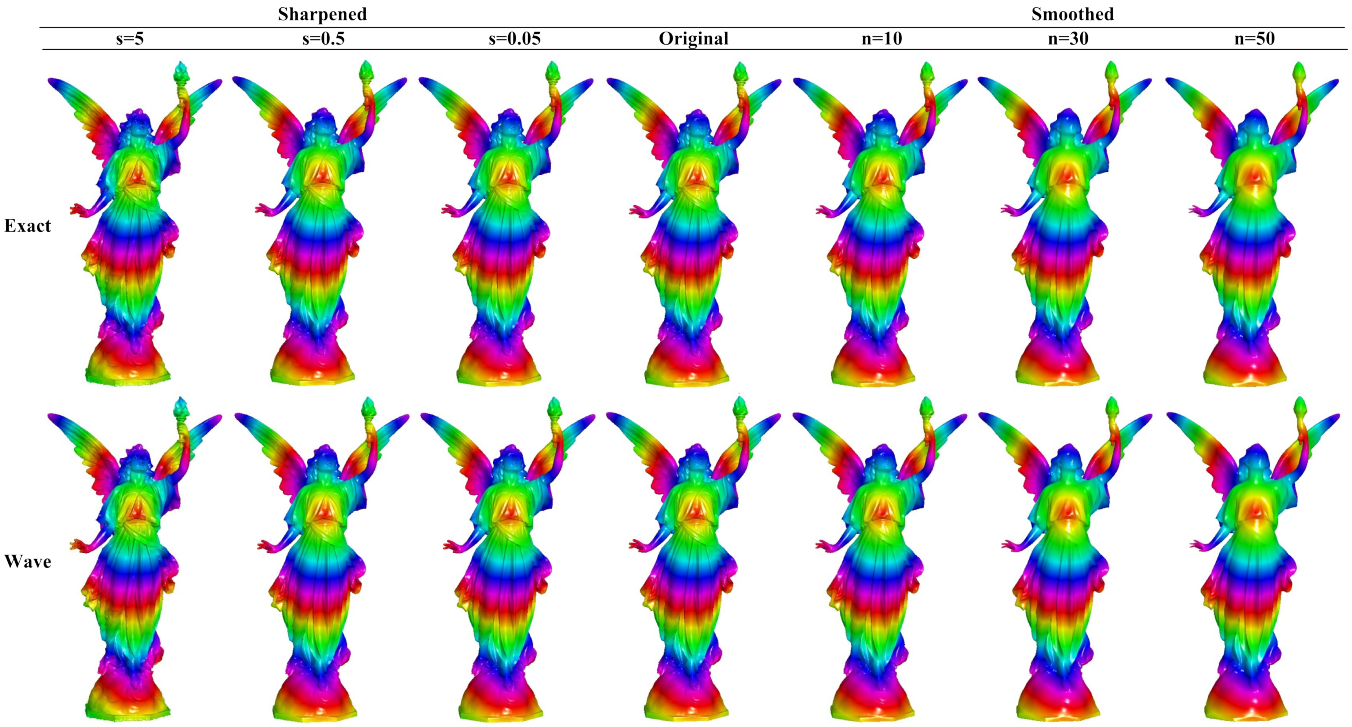

Figure 12: *Geodesic distances computed on the Lucy model. The top row shows the exact geodesics, while the bottom row shows geodesics from waves. From left to right, we have the Lucy model sharpened with* $s = 5, 0.5,$ *and* 0.05*, where* $\sigma = s * l$*, the original lucy model, and the lucy model smoothed with* $n = 10, 30,$ *and* 50 *iterations of simple umbrella smoothing.*

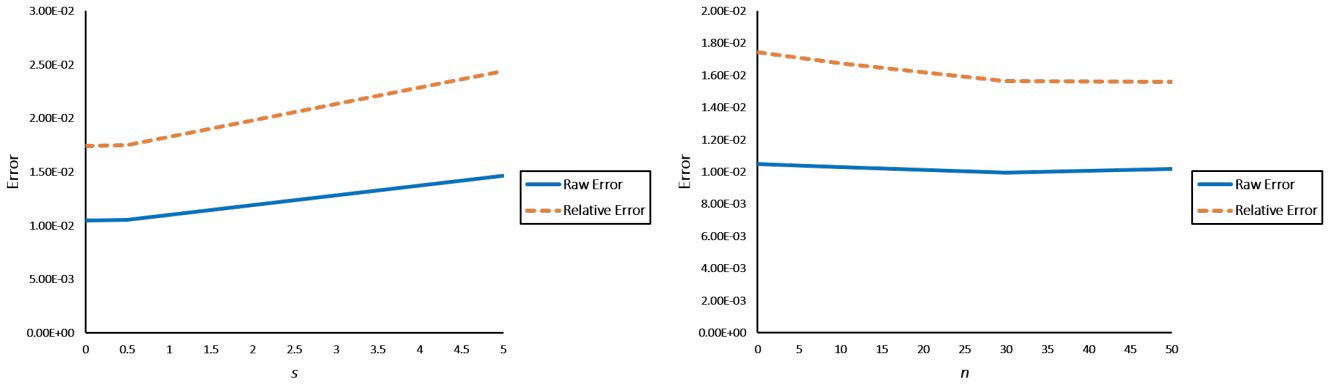

(a) **(b)**

Figure 13: *(a) Errors computed for increasing sharpness, and (b) errors computed for increasing smoothness.*

APPENDIX

A An alternative formulation for the computation of geodesics

In Eq. 9, we compute, for every edge, the mean of the gradients at the two triangles incident upon the edge, and save this mean gradient at the edge. This, however, leads to smoothing of gradients, and therefore, loss of information, which explains why the errors in Fig. 13a go up as the sharpness in the model increases.

Instead, we can save the normalized gradient, $\bar{G}$, at each triangle face, and compute the divergence of this gradient using the following equation:

$$Div_i = \int \langle \nabla h_i, \bar{G} \rangle,$$

where h_i is the hat-basis function at vertex v_i, and $\bar{G}$ is the normalized gradient per mesh face. Divergence can be discretized as:

$$Div_i(\bar{G}) = \sum_{f \in N(v_i)} \langle \nabla h_i \big|_f . \bar{G} \big|_f \rangle . |f|,$$

where the gradient of h_i is restricted to faces f that are incident upon the vertex v_i since the gradients of the hat-basis functions and the vector field, $\bar{G}$, are constant per face, and the gradient of the hat-basis functions is supported in the one-ring. $|f|$ is the area of triangle face f.

Similarly, the Laplacian for every vertex pair (i, j) can be computed as

$$L_{ij} = \int \langle \nabla h_i, \nabla h_j \rangle,$$

which can be discretized as

$$L_{ij} = \sum_{f \in N(v_i) \cap N(v_j)} \langle \nabla h_i \big|_f . \nabla h_j \big|_f \rangle |f|.$$

Since the Laplacian is the divergence of the gradient, we can solve the equation

$$\Delta w = Div_i(\bar{G})$$

to compute the underlying function, w, that approximates the geodesic distance function.

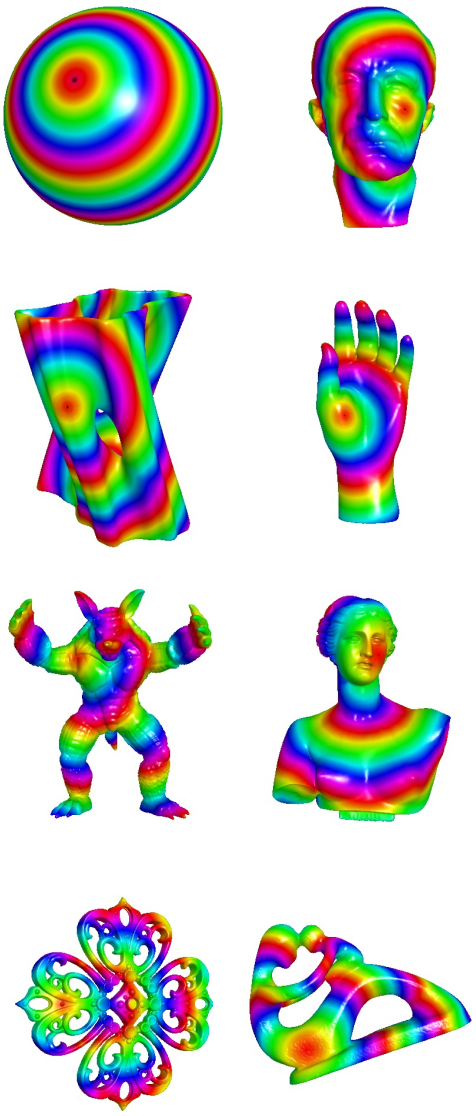

Figure 14: *Geodesics computed using our method on more models: (from left to right, top to bottom) Sphere, Max Planck, Julia vase, Hand, Armadillo man, Aphrodite, Filigree and Fertility models.*